Bridging Scales in Climate System Modeling

Todd Ringler

Los Alamos National Laboratory
Theoretical Division

LA-UR 09-00465

Climate, Ocean and Sea-Ice Modeling Project
http://public.lanl.gov/ringler/ringler.html

Los Alamos NATIONAL LABORATORY EST. 1943

Office of Science U.S. DEPARTMENT OF ENERGY

Acknowledgments:

Collaborators:

Max Gunzburger, Florida State University
Lili Ju, University of South Carolina
Joe Klemp, NCAR
Phil Jones, LANL
Jamal Mohd-Yosuf, LANL
Bill Skamarock, NCAR
John Thuburn, Exeter University

Research Topics:

Mesh Generation
Numerical Methods
Analysis of Large-Scale Dynamics
Analysis of Cloud-Resolving Scales of Motion
Implementation on hybrid CPU/GPU architectures

The Thesis of this Presentation:

Thesis: In the coming decade and beyond, the climate modeling community will be challenged to resolve scales and processes that are far beyond its current reach.

Support: There are currently unresolved processes that have a significant influence on the global climate system.

Support: There is a pressing need to quantitatively characterize the regional signature of anthropogenic climate change in order to assess the impacts on socio-economic activity.

Examples of processes that are currently unresolved.

Cloud Processes

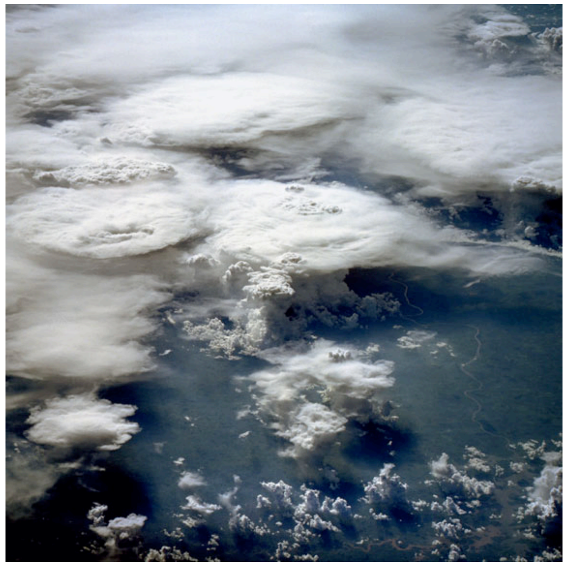

Ocean/Atmosphere Interaction

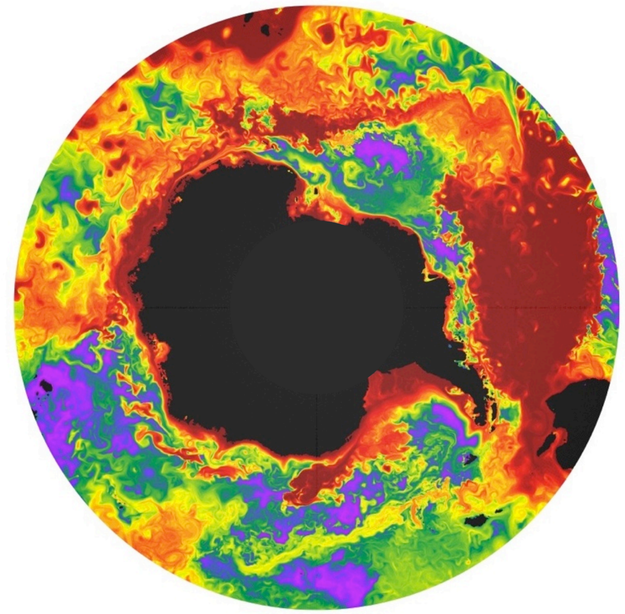

Ocean Biogeochemistry

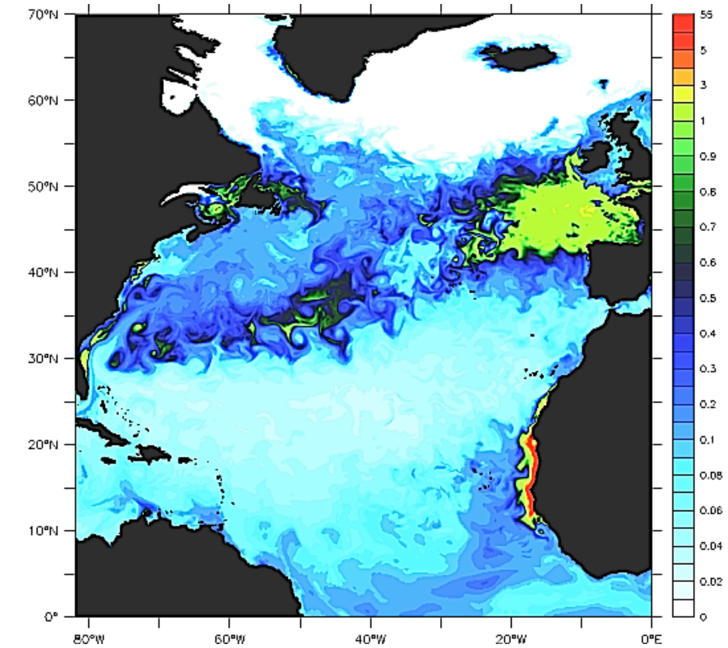

Each of these examples demonstrate scale-sensitive processes that might impact the climate system in a fundamental and important way.

The length scale of these processes is O(km).

Ocean/Ice Shelf Interaction

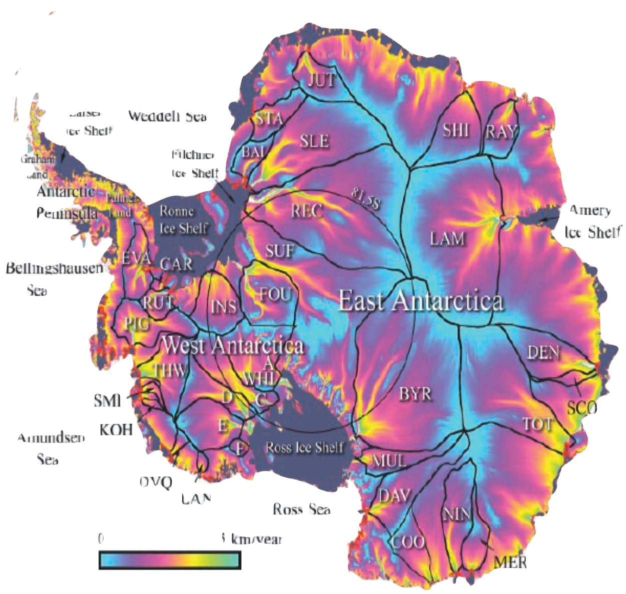

Hydrology in Complex Terrain

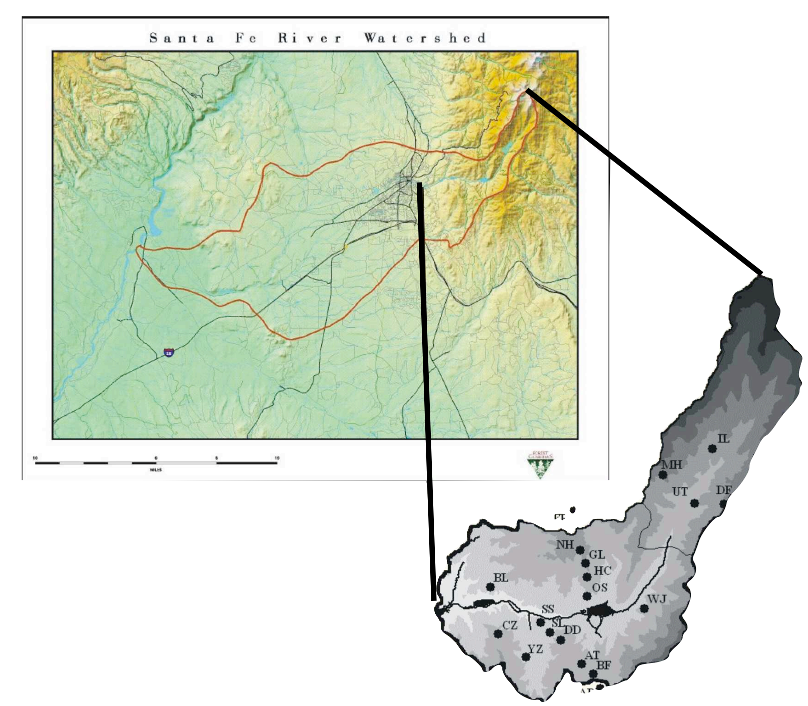

How do we get there?

A quick back-of-the-envelope analysis should be very concerning to mid-career scientists:

IPCC resolution ~ 100 km
Target resolution ~ 1 km*
Ratio of where we are to where we want to be ~ 2^7
Increase in computational resource required ~ 8^7 ~ 2e6
Time to reach target assuming a doubling in resources every 18 months ~ 22 years

Looking at the massive effort required to increase the resolution of our CCSMs, 22 years is not unrealistic. Maybe it is 15 years, or maybe it is 25 years. <u>Either way, it is a long time.</u>

The climate modeling community would benefit from an <u>additional</u> approach that allows <u>some</u> of these processes to be resolved at <u>some</u> locations.

We could then consider the notion of approaching the target from additional directions, for example by expanding the areas over which these fine-scale processes are simulated.

* While there is nothing magical about 1 km, it is a scale where some of the most troubling parameterizations are eliminated.

A Critique of the Global Quasi-Uniform Mesh Refinement Approach

Additional computing resources are distributed equitably across climate system component models, implying small increments in resolution (and thus in resolved processes) per upgrade.

As a result, we risk pushing the critical analysis of which small-scale processes are most important to the out-years.

Since we are always resource limited, reproducing the latest results requires the skills to use (and the access to) state-of-the-art high performance computing resources.

Since important scales are not resolved, we risk excluding important additional science components in our simulations of anthropogenic climate change.

With that said, the approach of global refinement has worked reasonably well for the last 30+ years. At this point in time there is no reason to think that it should not remain the primary vehicle for assessing climate change.

Still, there is room for a complementary approach that allows for the regional simulation of important climate processes ...... if we can figure out a way to do it.

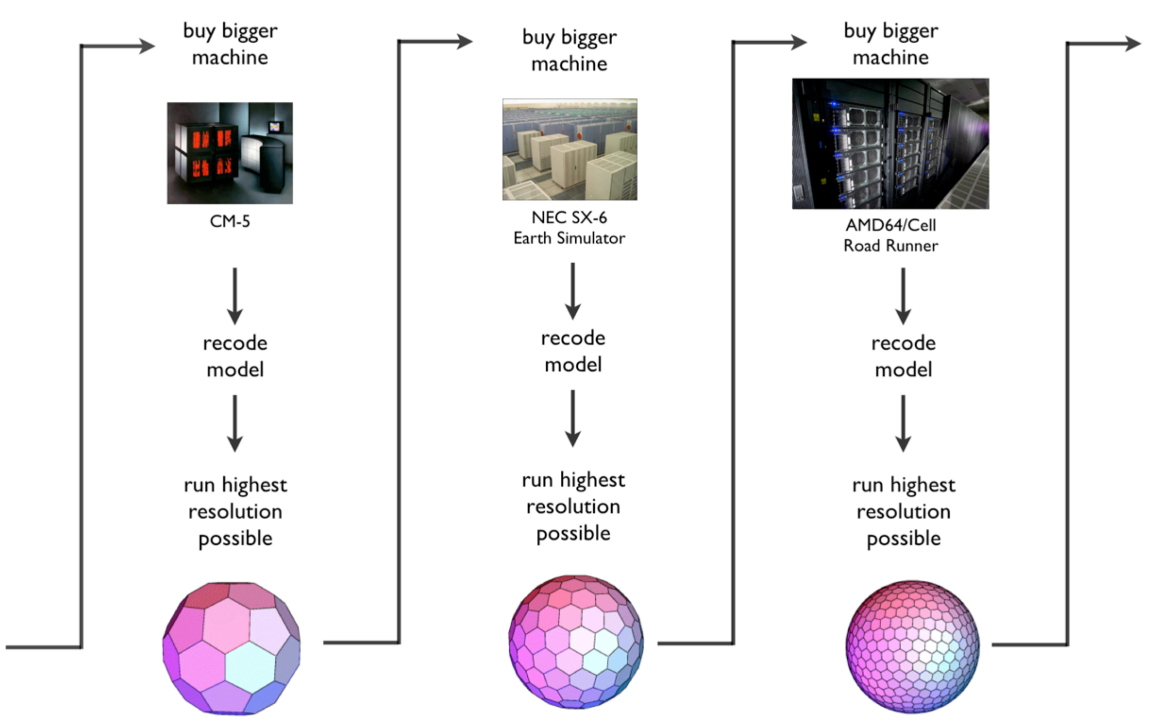

Roadblocks to a Multi-Scale Climate System Model

Roadblock #1: Numerical Error

Observation: IPCC-class climate system models are deeply-seated in the world of low-order (~2nd-order) finite-volume/finite-difference methods. Ocean models exclusively use these methods and atmosphere models are transitioning toward these methods.

Observation: Low-order methods are characterized by relatively large truncation errors than can quickly destroy simulations. One way to mitigate the problems of large truncation error is to use smooth grids. The implication here is that variable-resolution or nested meshes are problematic.

Observation: The bar-to-jump-over would appear to be fairly low here. Our primary motivation for using nested meshes is not the formal reduction in PDE solution error, but rather to permit the simulation of new phenomena.

Goal: Develop multi-scale finite-volume solvers that "do no harm."

Roadblocks to a Multi-Scale Climate System Model

Roadblock #2: Lack of Scale-Aware Physical Parameterizations

Observation: Increases in resolution are accompanied by a tremendous amount of effort to adjust sub-grid parameterizations. Examples include radiation parameterizations, cloud parameterizations, turbulence closures, etc. Changing global resolutions, even incrementally, is no small task.

Observation: To contemplate this adjustment process when a wide spectrum of resolutions is present in the same simulation is a difficult.

Goal: Develop a multi-scale approach that is tractable given the adjustment processes that must take place as resolution changes.

Roadblocks to a Multi-Scale Climate System Model

<u>Roadblock #1: Numerical Error</u>

<u>Roadblock #2: Lack of Scale-Aware Physical Parameterizations</u>

For a large class of meshes, we have lowered (and possibly even removed) Roadblock #1 while making choices that mitigate Roadblock #2.

The demonstration will go as follows:

1. Discussion of nested conformal, variable-resolution meshes
2. Constraining truncation error on these meshes
3. A demonstration via numerical simulations

Conformal, Variable-Resolution Meshes

A conformal mesh is a mesh with no hanging nodes.

conformal

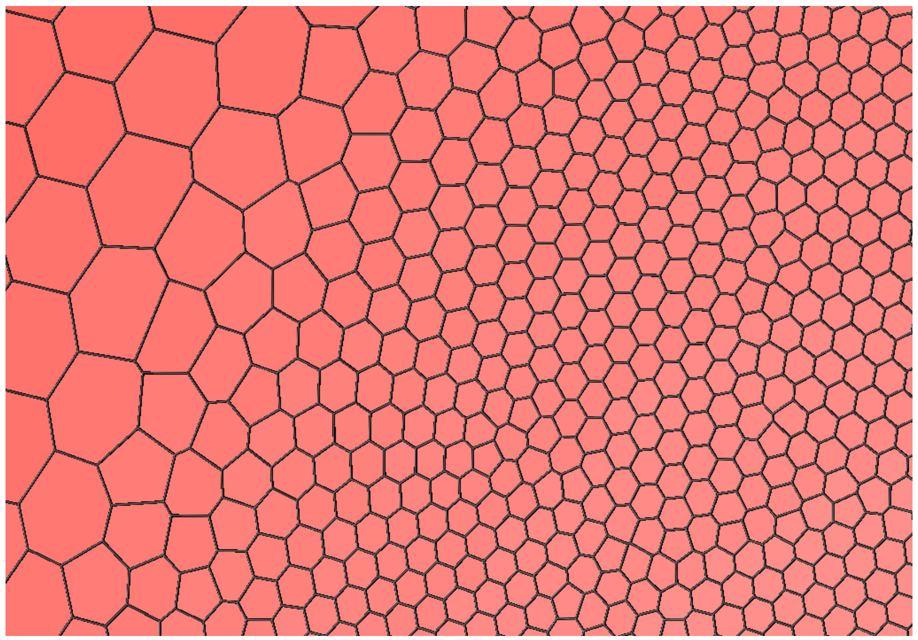

non-conformal

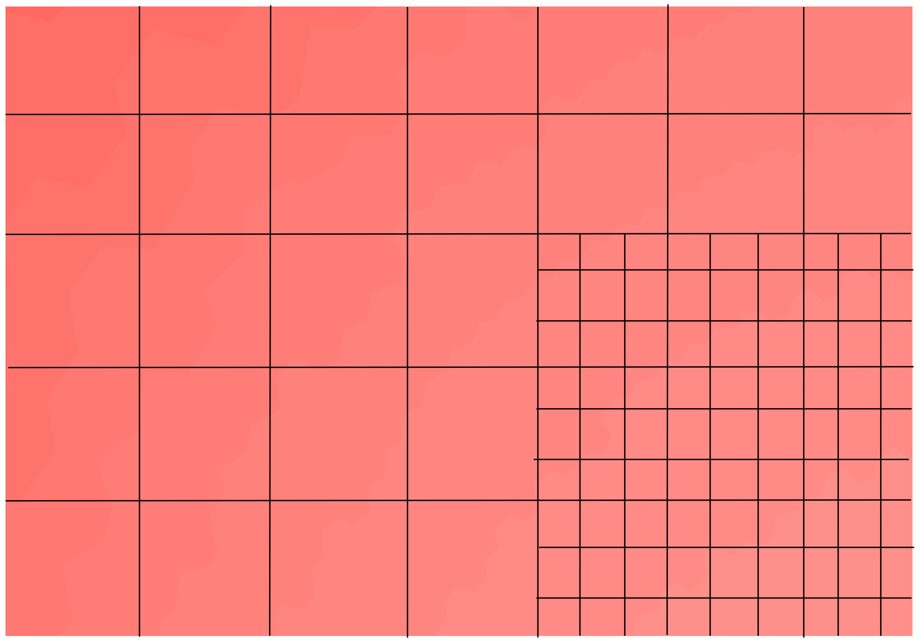

We are developing algorithms that target conformal meshes. These meshes include arbitrary Voronoi tessellations and Delaunay triangulations (as well as the standard Lat/Lon and conformally-mapped cubed sphere.)

Nested, Conformal, Variable-Resolution Meshes

In order to mitigate our lack of scale-aware physical parameterizations, we are focused on a nested approach.

The approach is equally valid for atmosphere, ocean, and ice modeling.

This has the added benefit of allowing multiple time-steps. A short time step in the nested domain and a longer time step in the coarse domain.

We are essentially taking a page out of the numerical-weather-prediction playbook.

Nested, Conformal, Variable-Resolution Meshes

In order to mitigate our lack of scale-aware physical parameterizations, we are focused on a nested approach.

The approach is equally valid for atmosphere, ocean, and ice modeling.

This has the added benefit of allowing multiple time-steps. A short time step in the nested domain and a longer time step in the coarse domain.

We are essentially taking a page out of the numerical-weather-prediction playbook.

Nested, Conformal, Variable-Resolution Meshes

In order to mitigate our lack of scale-aware physical parameterizations, we are focused on a nested approach.

The approach is equally valid for atmosphere, ocean, and ice modeling.

This has the added benefit of allowing multiple time-steps. A short time step in the nested domain and a longer time step in the coarse domain.

We are essentially taking a page out of the numerical-weather-prediction playbook.

Nested, Conformal, Variable-Resolution Meshes

In order to mitigate our lack of scale-aware physical parameterizations, we are focused on a nested approach.

The approach is equally valid for atmosphere, ocean, and ice modeling.

This has the added benefit of allowing multiple time-steps. A short time step in the nested domain and a longer time step in the coarse domain.

We are essentially taking a page out of the numerical-weather-prediction playbook.

Nested, Conformal, Variable-Resolution Meshes

In order to mitigate our lack of scale-aware physical parameterizations, we are focused on a nested approach.

The approach is equally valid for atmosphere, ocean, and ice modeling.

This has the added benefit of allowing multiple time-steps. A short time step in the nested domain and a longer time step in the coarse domain.

We are essentially taking a page out of the numerical-weather-prediction playbook.

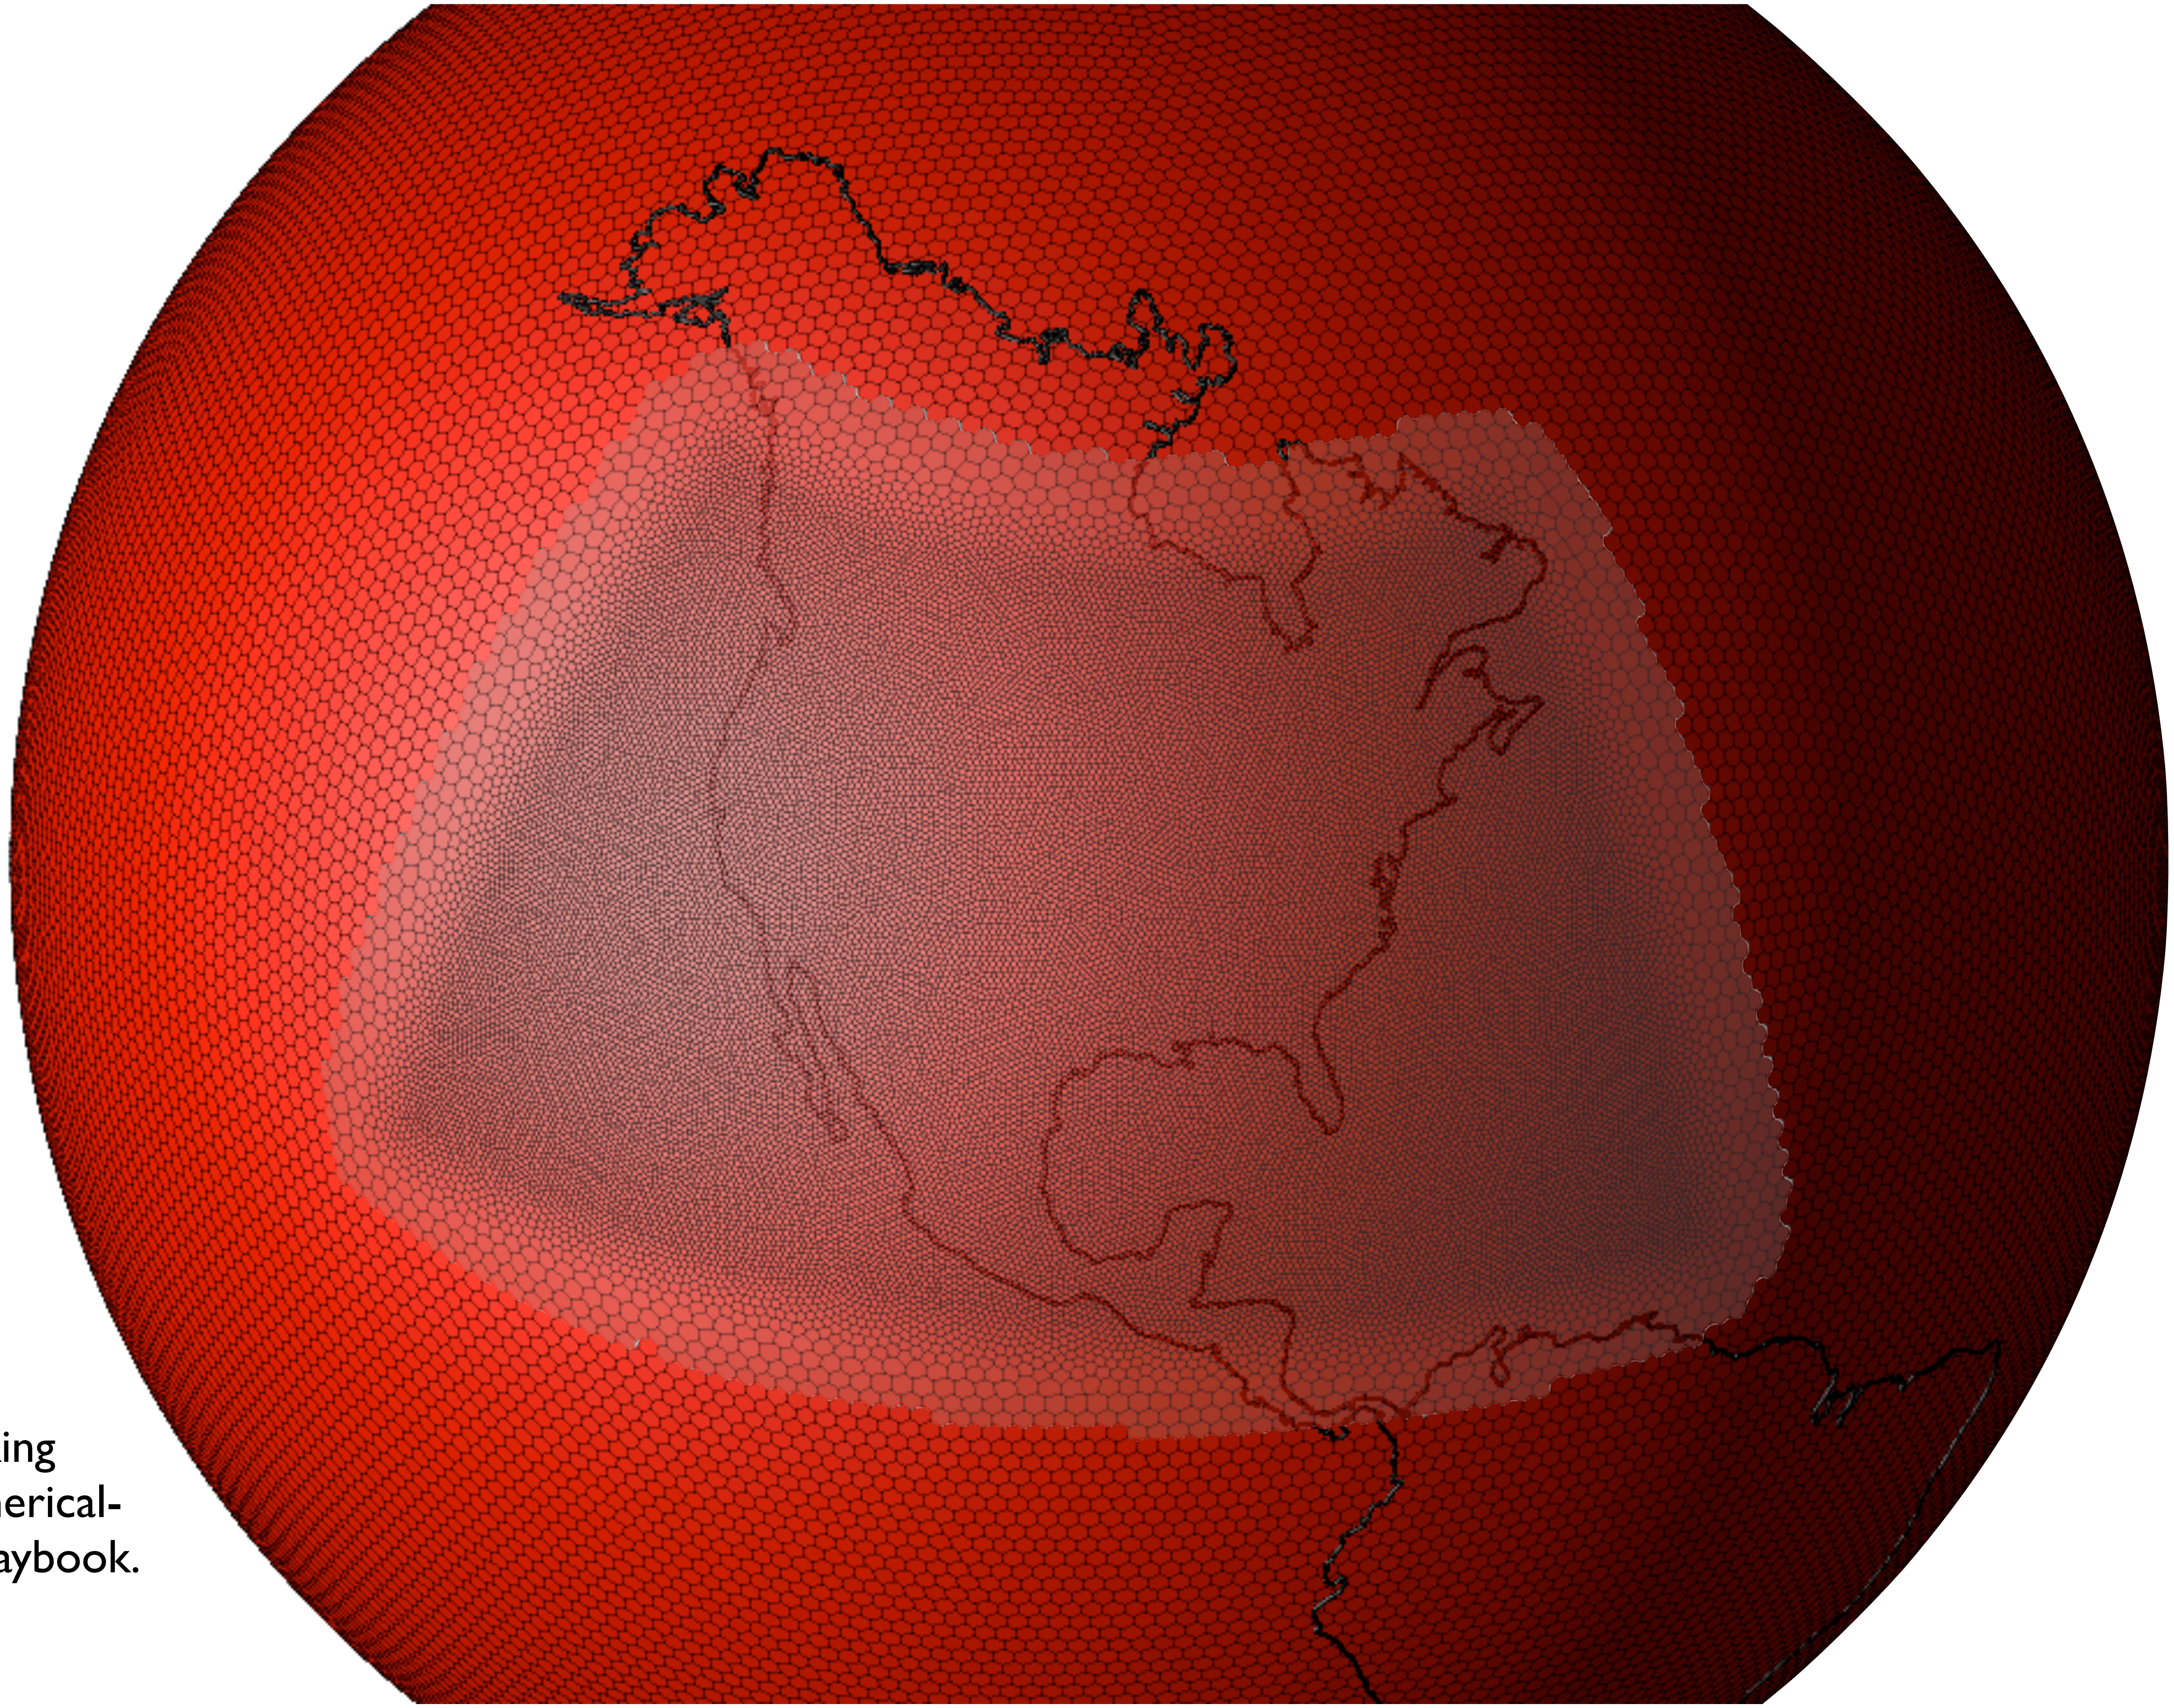

A mesh does not a simulation make.

Significant advances in the construction of numerical algorithms on C-grids have been obtained.

Analytic Results:

1. Stationary geostrophic mode is recovered.
2. Total energy is conserved to within time truncation.
    a. Coriolis force is energetically-neutral
    b. Transport of KE is conservative
    c. KE/PE exchange is equal and opposite.
3. Potential vorticity is conserved to round-off.
4. It appears* that potential enstrophy can be dissipated.

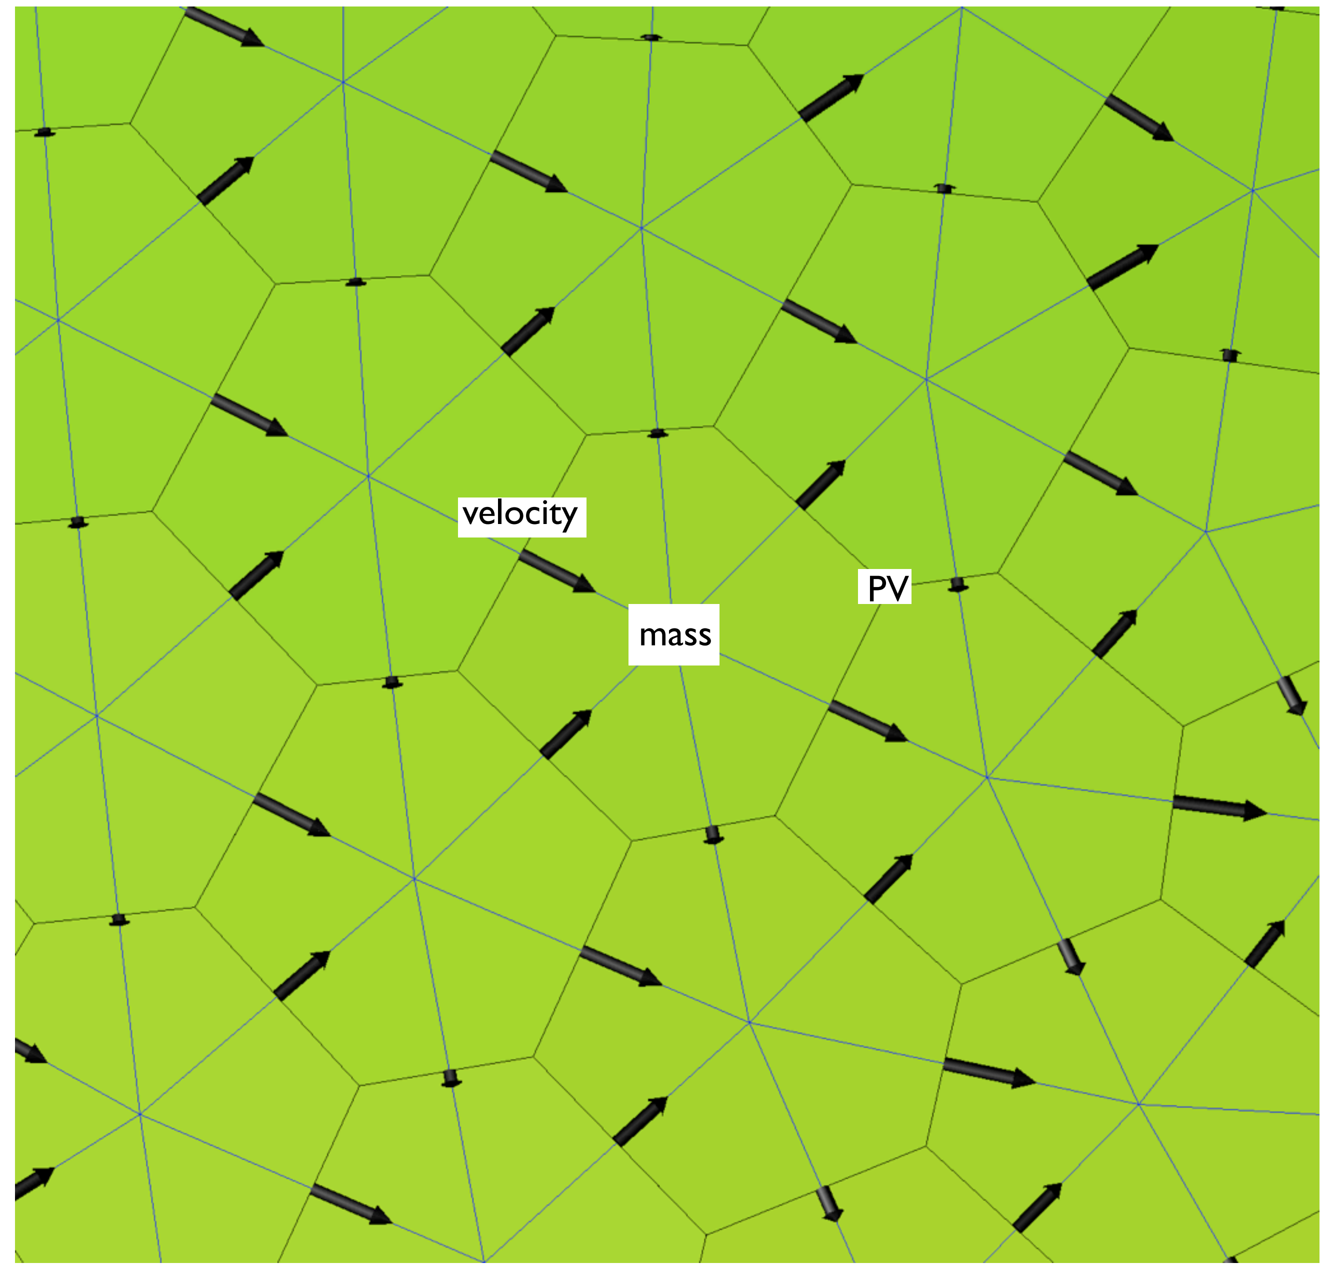

Results hold for a wide class of meshes: Lat/Lon, Voronoi Tessellations, Delaunay Triangulation and Conformally-mapped cubed sphere meshes.

*appears likely, but not yet proven.

Thuburn, J., T. Ringler, W. Skamarock and J. Klemp, 2009: Numerical treatment of geostrophic modes on arbitrarily structured C-grids. JCP, in preparation.

Ringler, T., J. Thuburn, W. Skamarock and J. Klemp, 2009: Numerical treatment of energy and potential vorticity on arbitrarily structured C-grids. JCP, in preparation.

Ringler, T., L. Ju, M. Gunzburger, J. Thuburn, W. Skamarock and J. Klemp, 2009: Simulation of the shallow-water equations on variable-resolution C-grids. MWR, in preparation.

Numerical Demonstration using Shallow-Water Equations

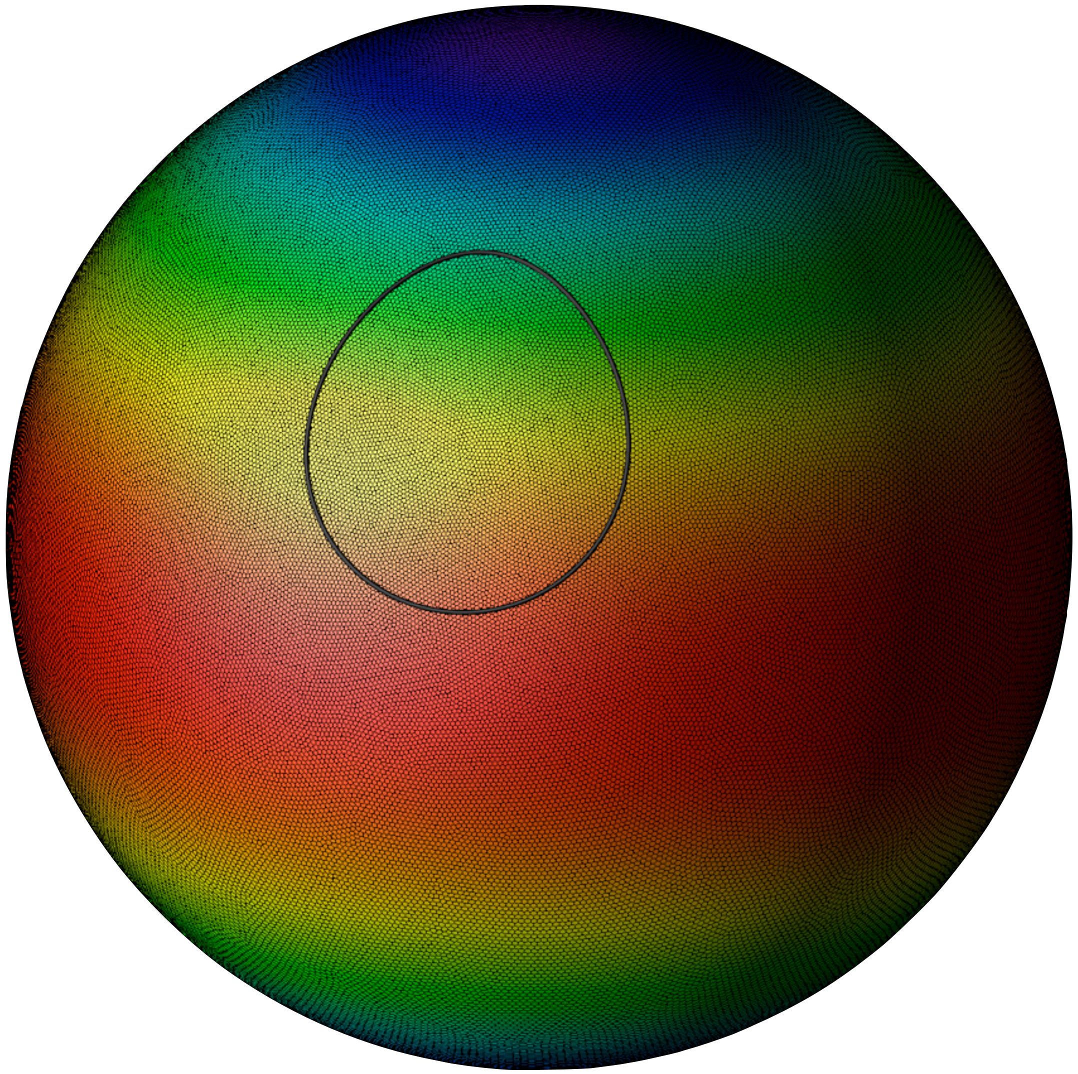

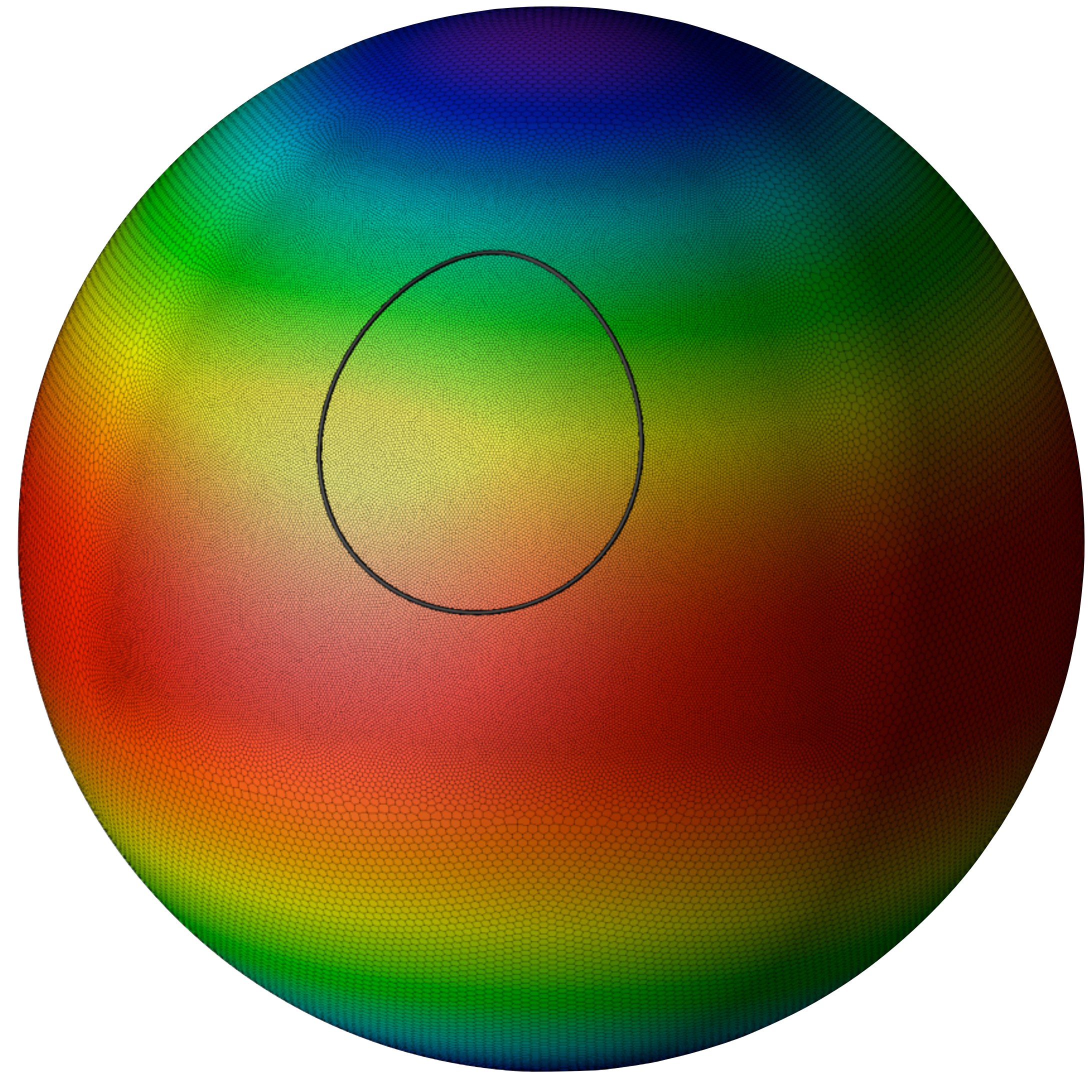

82547 nodes distributed
in a quasi-uniform manner.

82547 nodes with focus
in region of orography.

<u>Simulation details</u>: RK4 time integration, centered-in-space numerics, no dissipation.
15 day integration. (Much longer integrations with no dissipation are possible.)

Results: Day 15, thickness field (m)

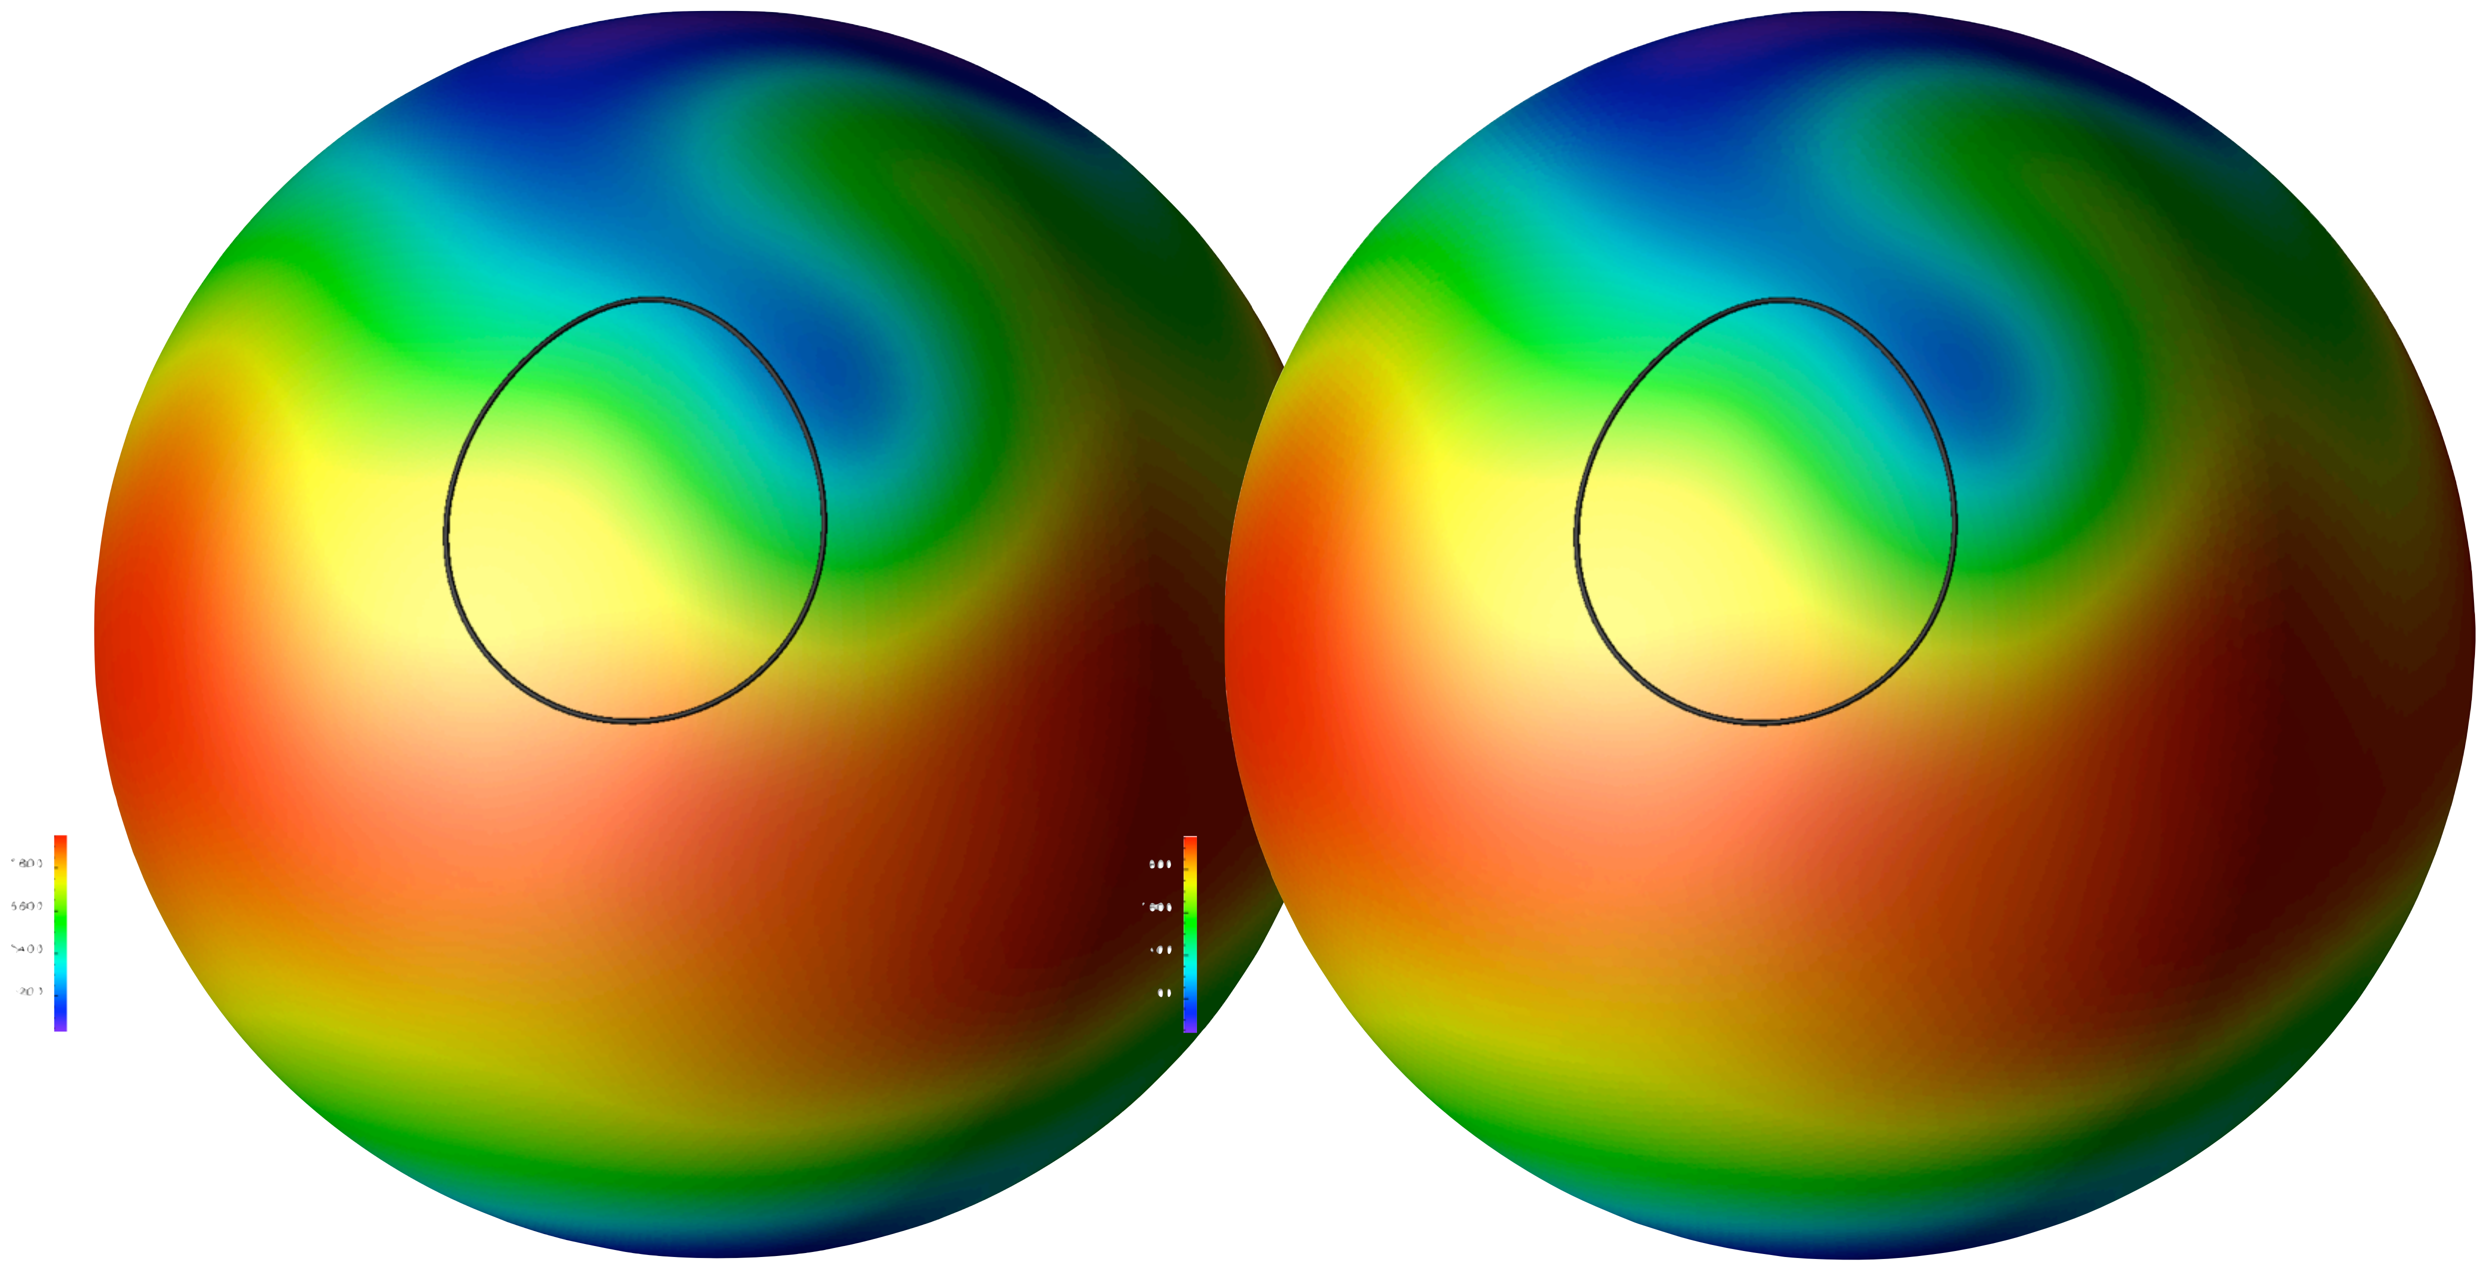

Indistinguishable.

Results: Day 15, relative vorticity (1/s)

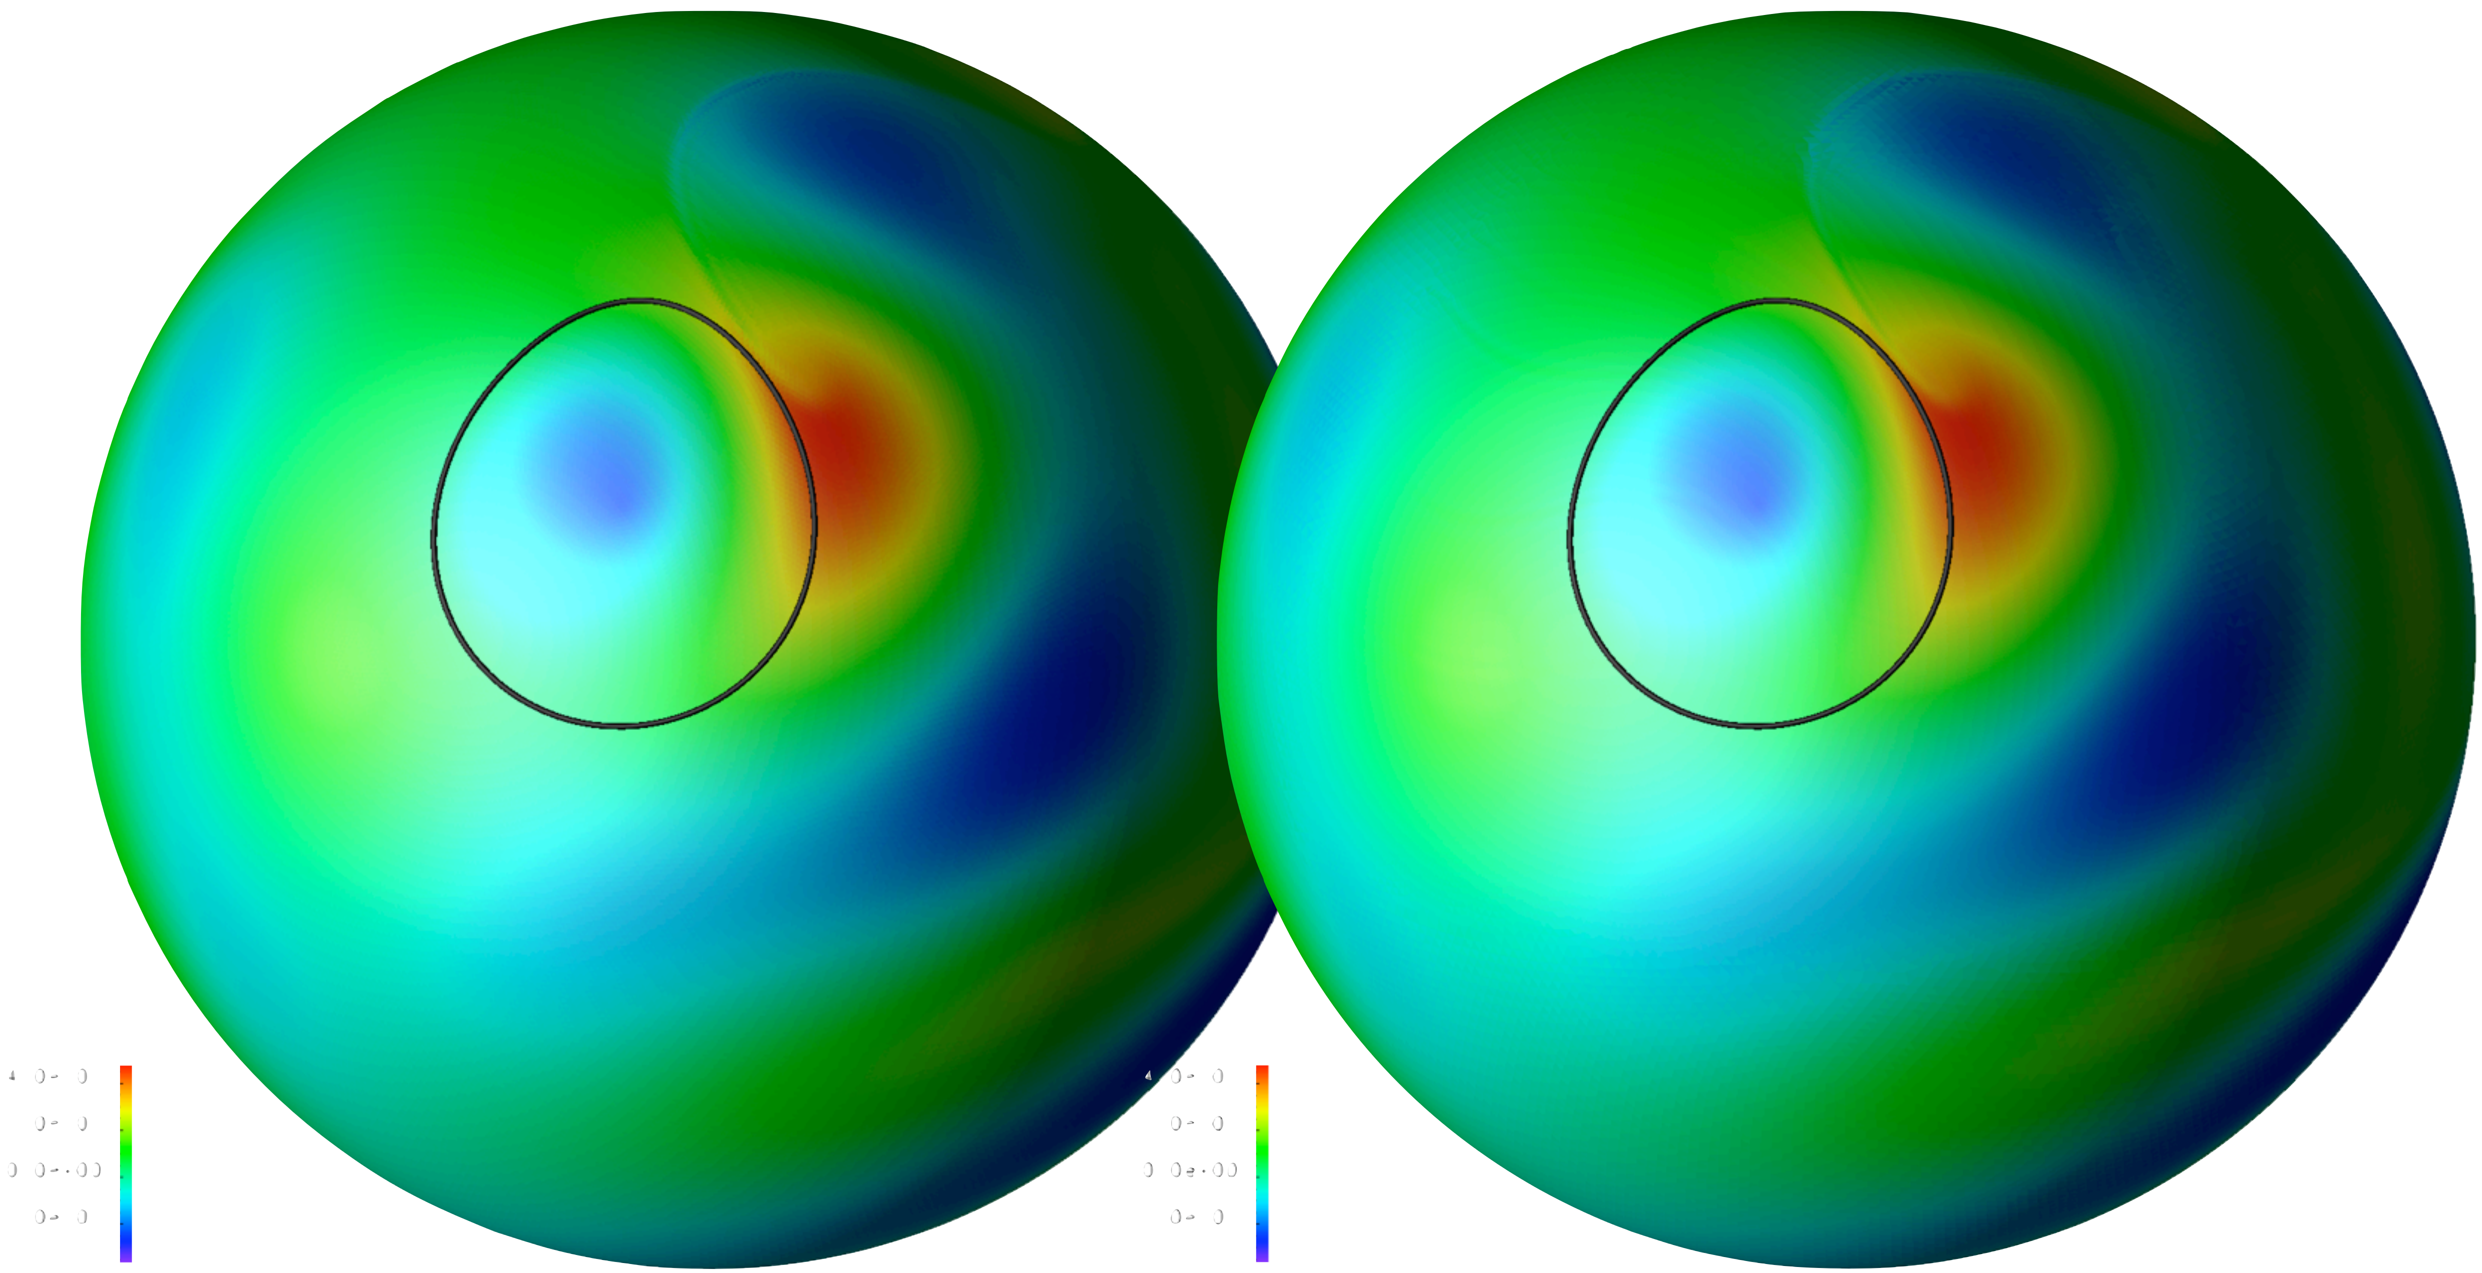

Essentially Indistinguishable.

Results: Error Norms of thickness at day 15

Errors computed with respect to a T216 spectral solution.

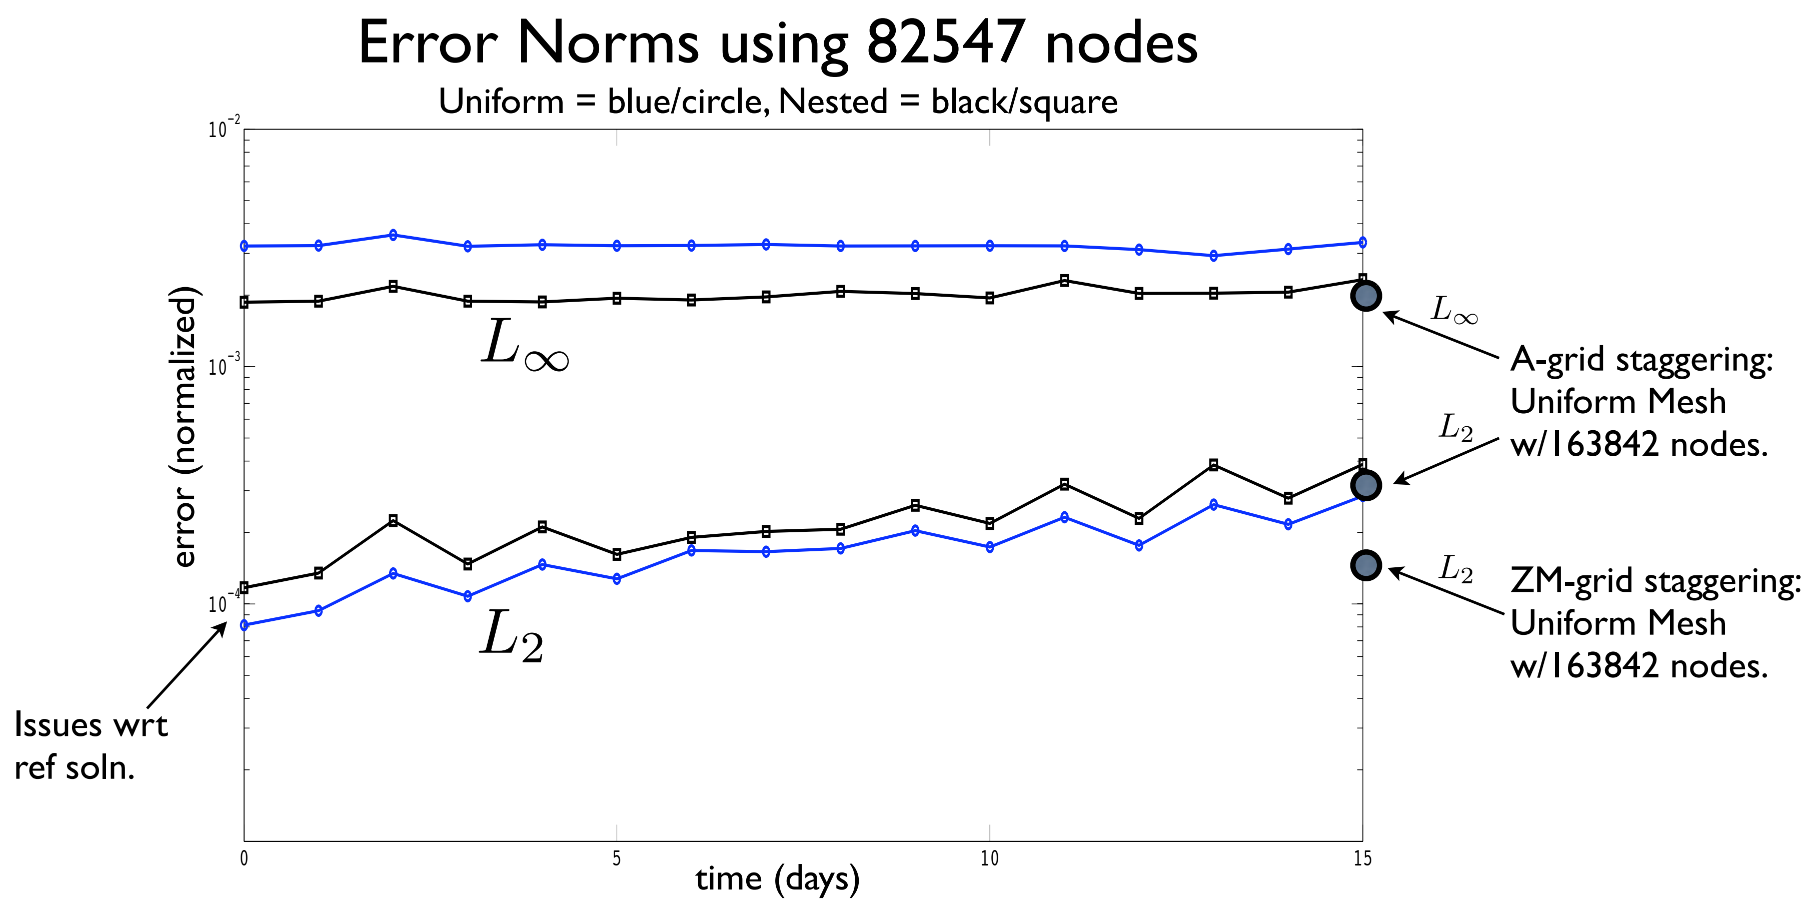

82547 nodes, uniform

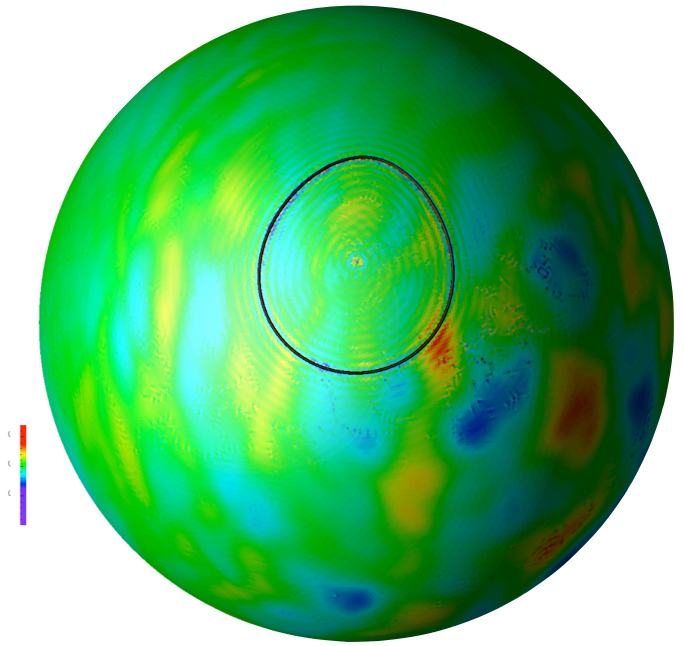

82547 nodes, nested

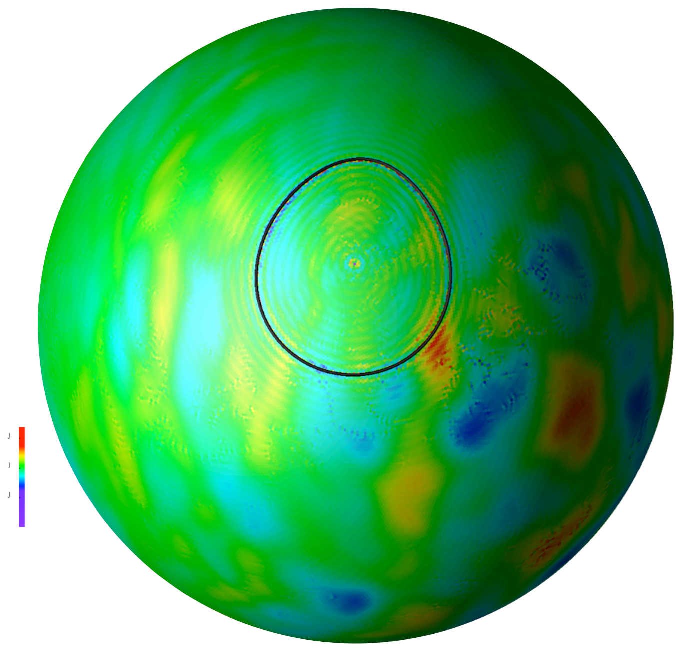

Errors are plotted on the same scale.

1. Simulation with nested mesh is comparable to simulation with uniform mesh.
2. Both simulations compare favorably with simulations using other grid staggerings.
3. For this simulation, we have jumped over the do-no-harm bar.

3-D Supercell Simulation ~2000 m Horizontal Grid
(Effort lead by Klemp and Skamarock.)

This numerical approach performs well at cloud-resolving scales.

rectangular mesh

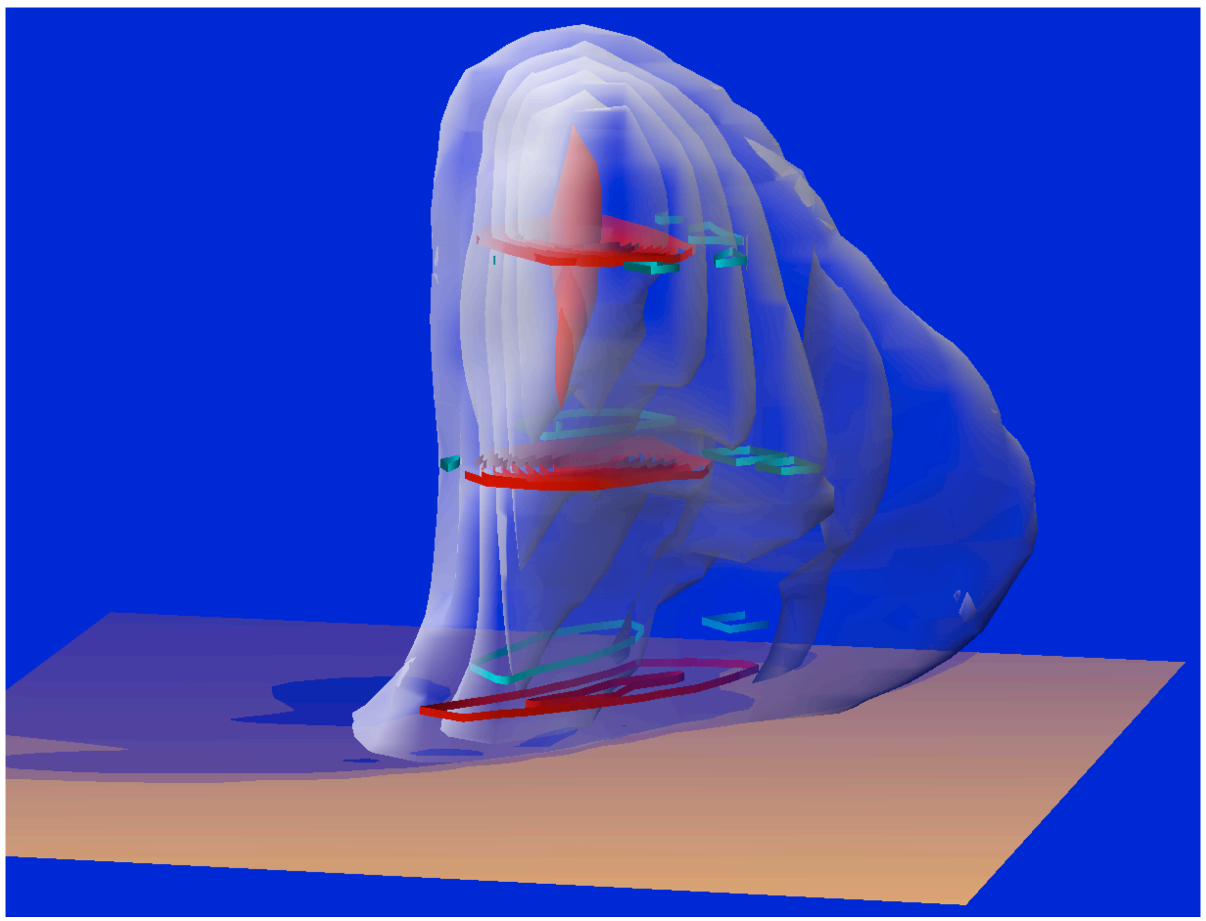

Voronoi tessellation (hexagons)

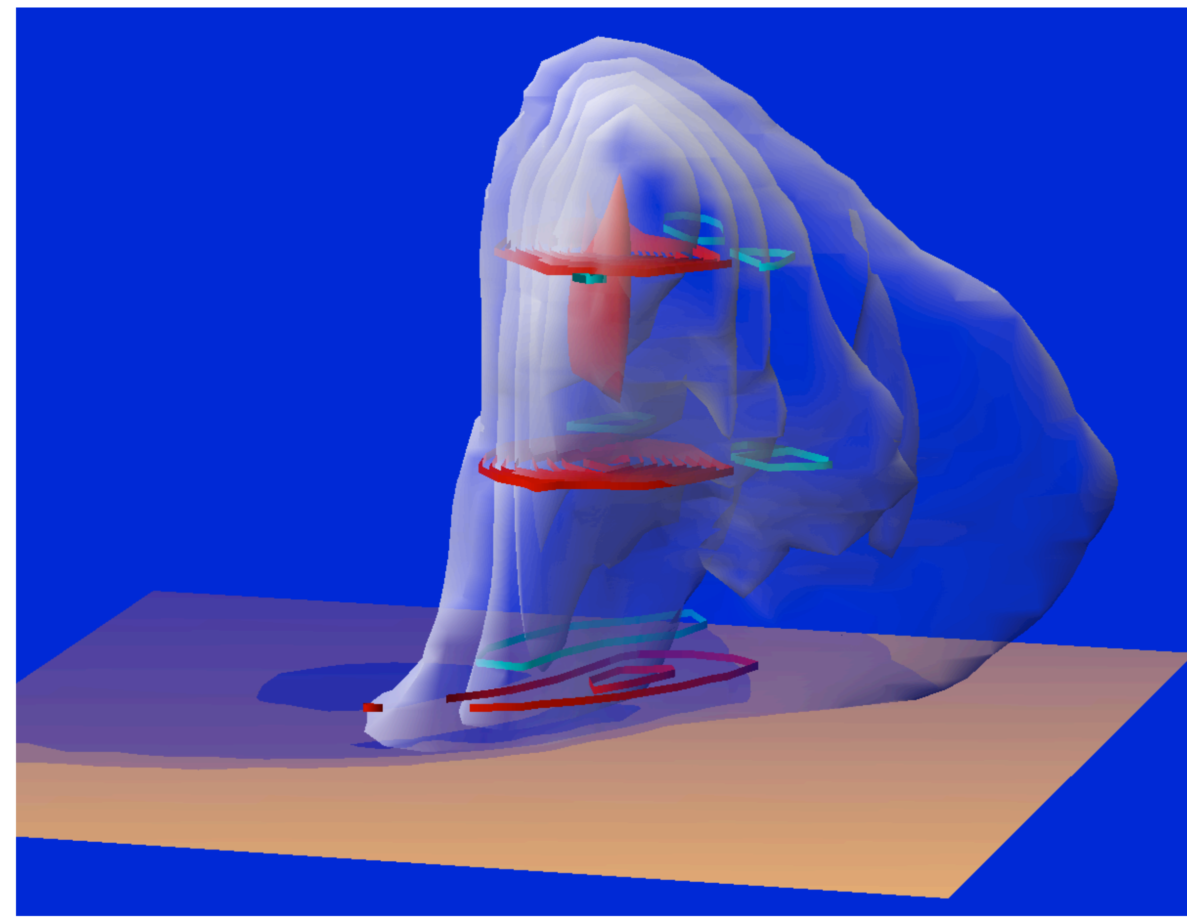

vertical velocity contours at 1, 5 and 10 km (c.l = 3 m/s)
30 m/s vertical velocity surfaces shaded in red
rainwater surfaces shaded as transparent shells
perturbation surface temperature shaded on baseplane.

We think that this approach is analytically sound .... but can we get the throughput?

This is joint LANL / NCAR MMM development project.

Approach:
The method requires an unstructured grid approach.
Targeting hybrid CPU/GPU machines first.
Start with a stacked shallow-water model.
Vectorize over the vertical layer index.
Group operators together and push to GPU.

Assuming that we can get the required CPU efficiency and scaling, LANL will likely pursue the construction of an ocean dynamical core and NCAR MMM will consider the construction of an atmosphere dynamical core.

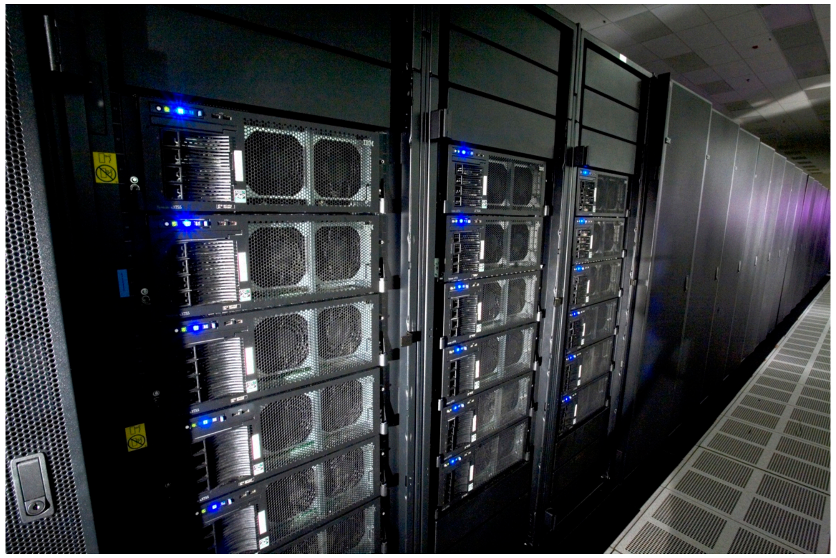

LANL Cerrillos

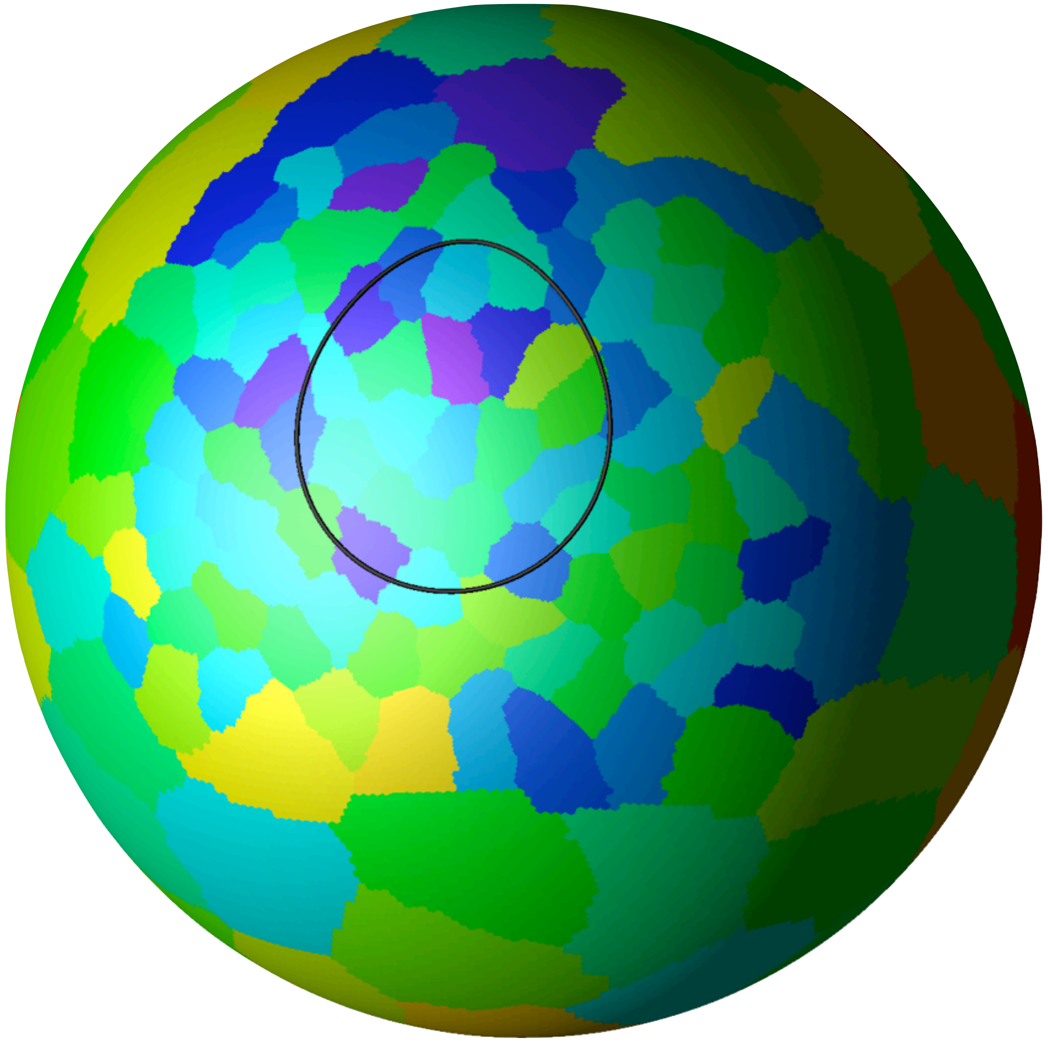

Block decomposition:
Distribute blocks across processors
Maximize area to circumference ratio
Use for load balancing.

Summary

The potential scientific-payoff in developing a climate system model with the ability to provide regional focus is tremendous.

The approach outlined here argues for regional climate process modeling within the framework of a global climate system model.

The ability to provide regional focus within a global climate system model would be complementary to the standard IPCC approach.

By exploring the importance of resolving fine-scale phenomena in a model with regional focus, we can provide insights into how the global quasi-uniform climate system model should evolve in time.

We have diminished (or even removed) one of the two primary roadblocks to building a multi-scale climate system model based on finite-volume methods.

Assuming that we can develop computational-efficient algorithms for this approach, we will continue to explore the potential merits of this approach.

Modus Operandi of Climate System Modeling

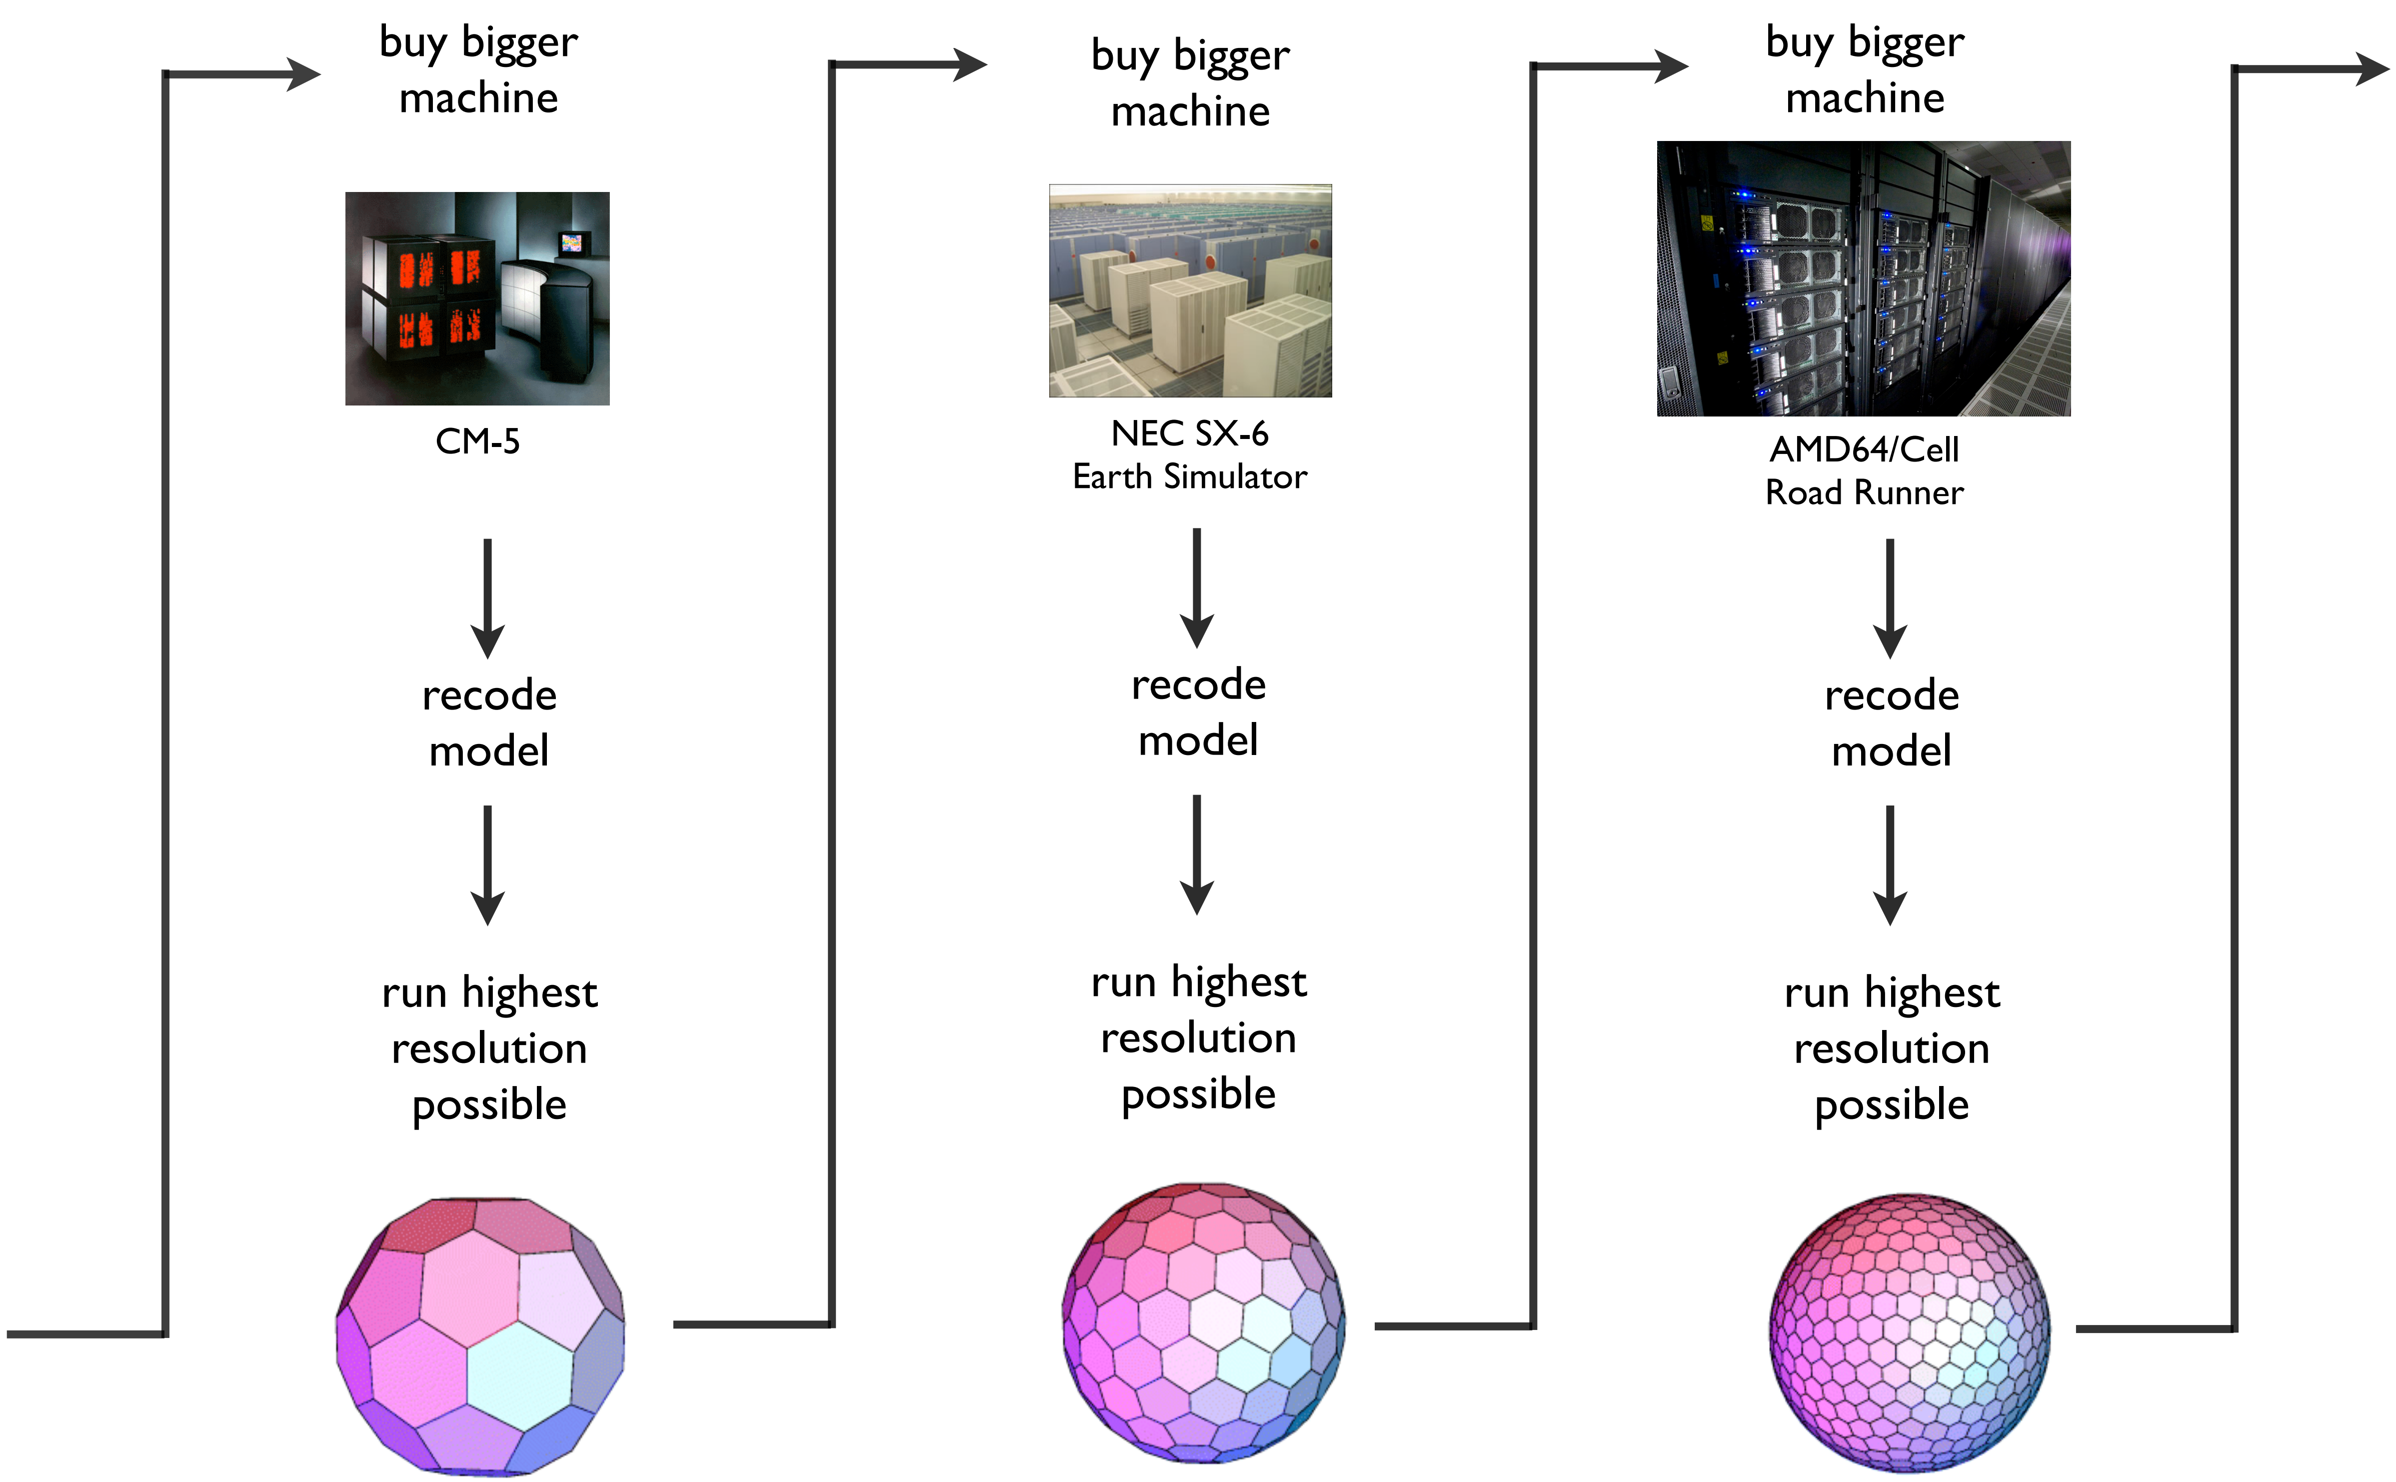

There is absolutely nothing wrong with this approach. Clearly it works.
The fact that it works should not preclude other complementary approaches.